

ABSTRACT

In an example data mining process, performance data for instructions that
5 execute in a data processing system is obtained. The performance data may
comprise instruction addresses and corresponding performance information. A dump
that comprises the instructions and corresponding instruction addresses may also be
obtained. Common code segments in the dump may be automatically identified. A
common code segment may comprise an ordered set of multiple instructions that
10 appears multiple times in the dump. Aggregate performance data for the common
code segments may be generated, based at least in part on (a) the instruction
addresses associated with the common code segments in the dump, and (b) the
instruction addresses and the corresponding performance information from the
performance data. Other embodiments are described and claimed.

15